CISC 7332X T6

C14b: Classless Intradomain Routing

Hui Chen

Department of Computer & Information Science

CUNY Brooklyn College

Acknowledgements

- Some pictures used in this presentation were obtained from the Internet
- The instructor used the following references
  - Larry L. Peterson and Bruce S. Davie, Computer Networks: A Systems Approach, 5th Edition, Elsevier, 2011
  - Andrew S. Tanenbaum, Computer Networks, 5th Edition, Prentice-Hall, 2010
  - James F. Kurose and Keith W. Ross, Computer Networking: A Top-Down Approach, 5th Ed., Addison Wesley, 2009
  - Larry L. Peterson's (http://www.cs.princeton.edu/~llp/) Computer Networks class web site

Outline

- Problem to scale to global network
  - Many networks organized in hierarchical manner
  - Scarcity of IP address
- Solution
  - Subnetting (legacy)
  - Supernetting (classless routing, legacy)
  - Classless routing (CIDR)

Is 2^{32} too small a number?

- IPv4 address
  - 32 bit integers ~ 2^{32} = 4,294,967,296
  - Many addresses, but not too many networks!

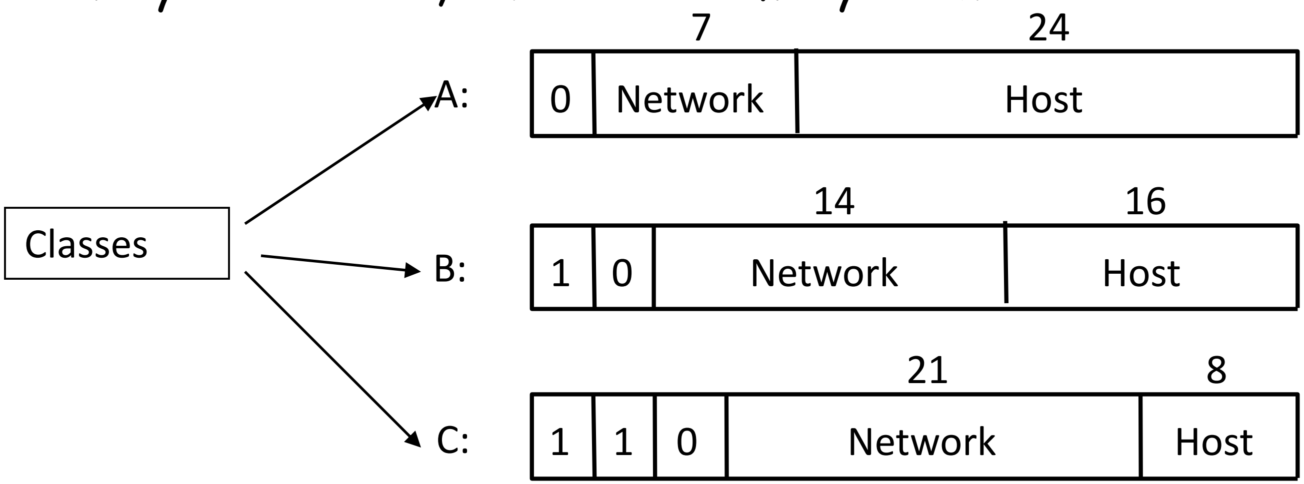

  - Testimony: http://www.iana.org/assignments/ipv4-address-space/
  - Examples
    - A network of two nodes needs a class C network
    - A network of 256 nodes needs a class B network

Can the number of networks be too many?

- How many class B networks are there?
- Potentially how big a routing table can be?

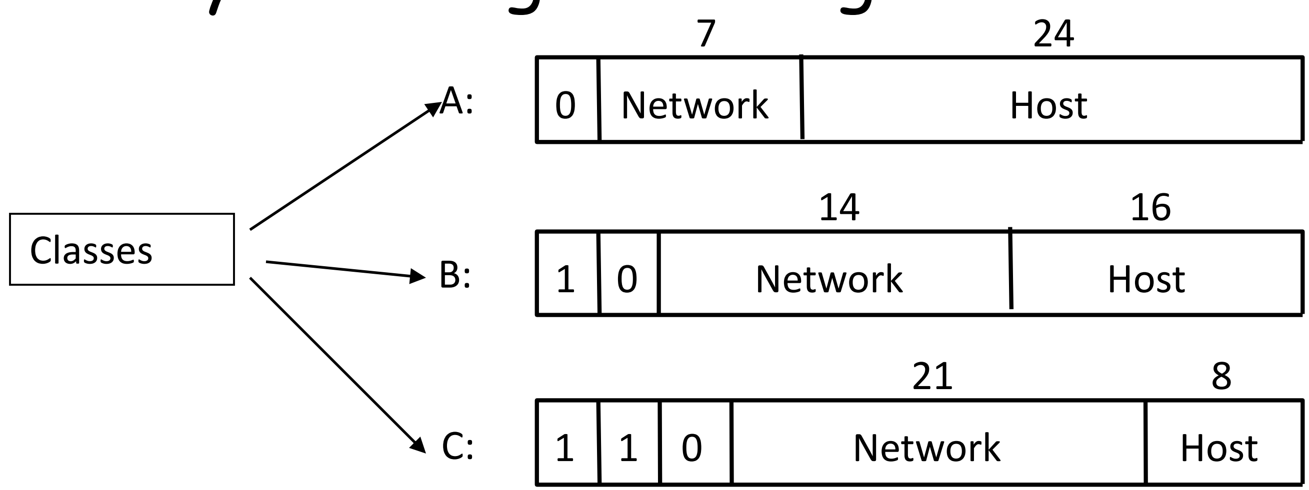

Subnetting

- Add another level to address/routing hierarchy: *subnet*
- *Subnet masks* define variable partition of host part of class *A* and B addresses
- Subnets visible only within site

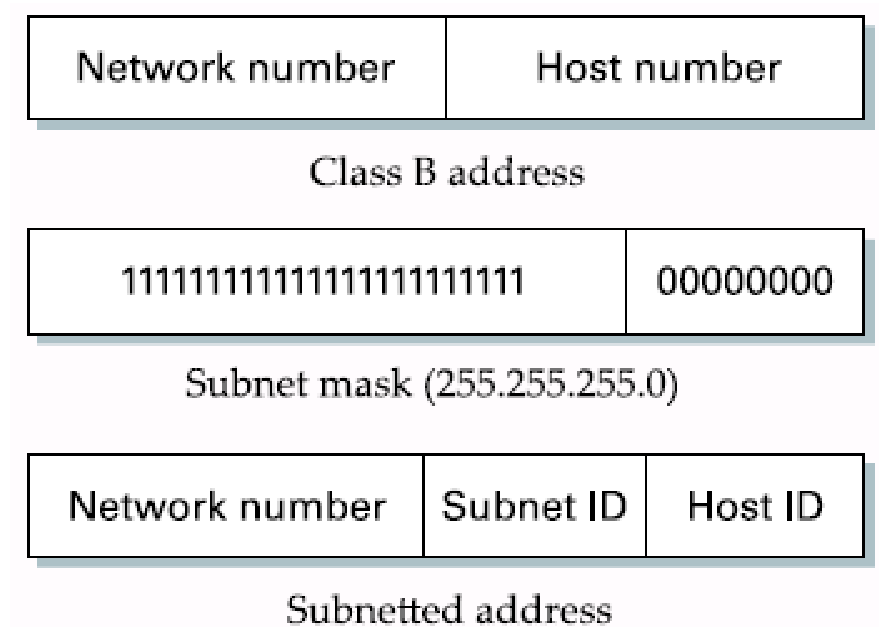

Subnetting: Example

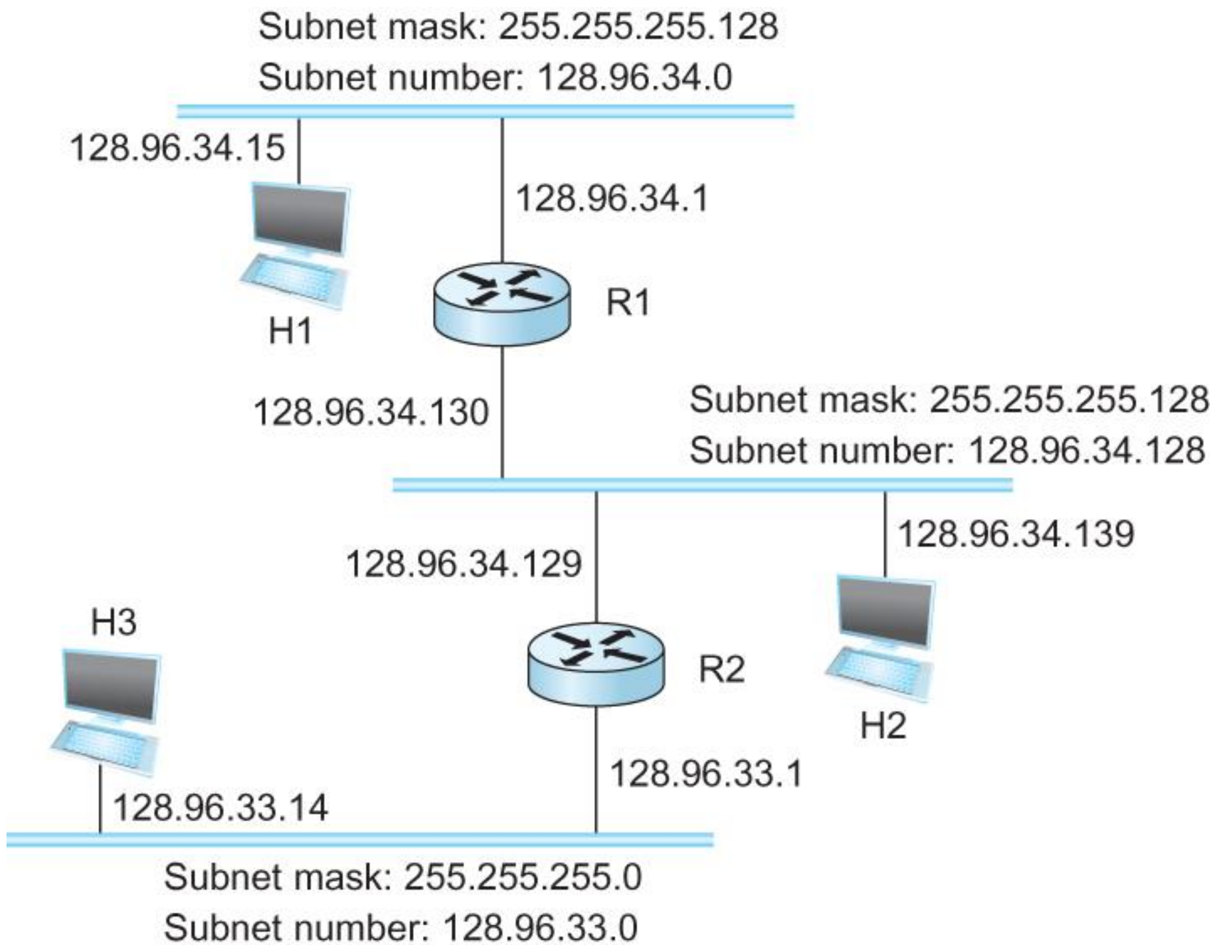

Subnetting: Example

- Forwarding Table at Router R1

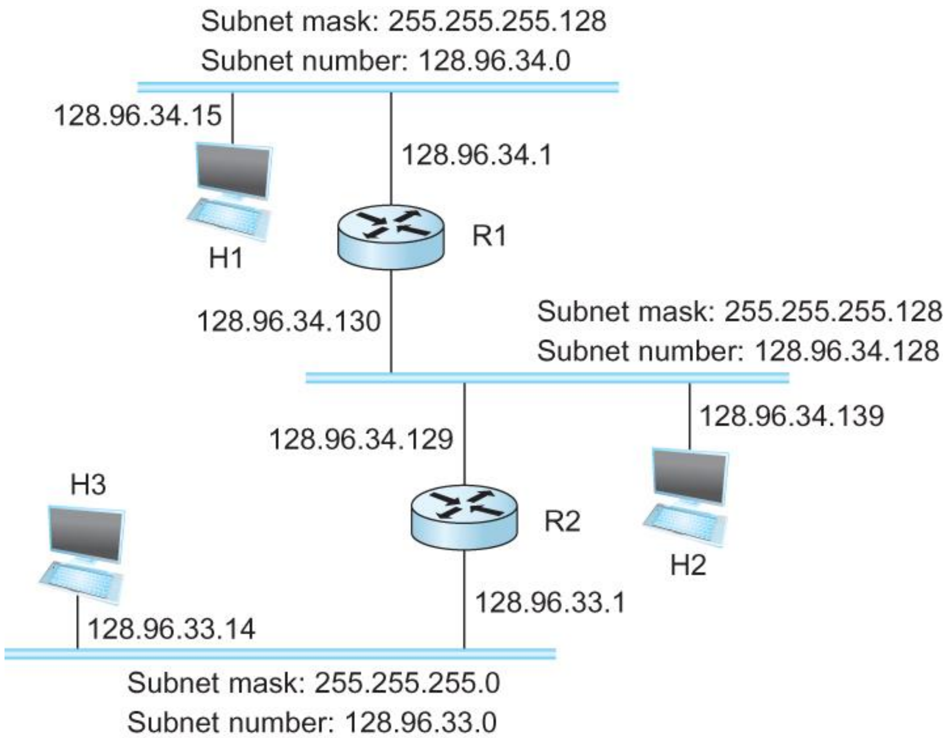

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Forwarding Algorithm

```
D = destination IP address
for each entry < SubnetNum, SubnetMask, NextHop>
   D1 = SubnetMask & D
   if D1 = SubnetNum
      if NextHop is an interface
         deliver datagram directly to destination
      else
         deliver datagram to NextHop (a router)
```

Subnetting: Discussion

- Would use a default router if nothing matches
- Subnet masks do not have to align with a byte boundary
- Subnet masks need **not** to be contiguous 1's
  - 255.255.1.0 is OK
    - 11111111 11111111 00000001 00000000
    - What is subnet number of IP address 128.96.34.1?

      10000000 01100000 00100010 00000000 &

      11111111 11111111 00000001 00000000 →

      10000000 01100000 00000000 00000000 →

      128.96.0.0 → can not directly tell from the IP address
  - In practice, use contiguous 1's
- Multiple subnets can be on a single physical network
- Subnets not visible from the rest of the Internet

Subnetting: Discussion

- How do you tell whether an IP address is on a given subnet?

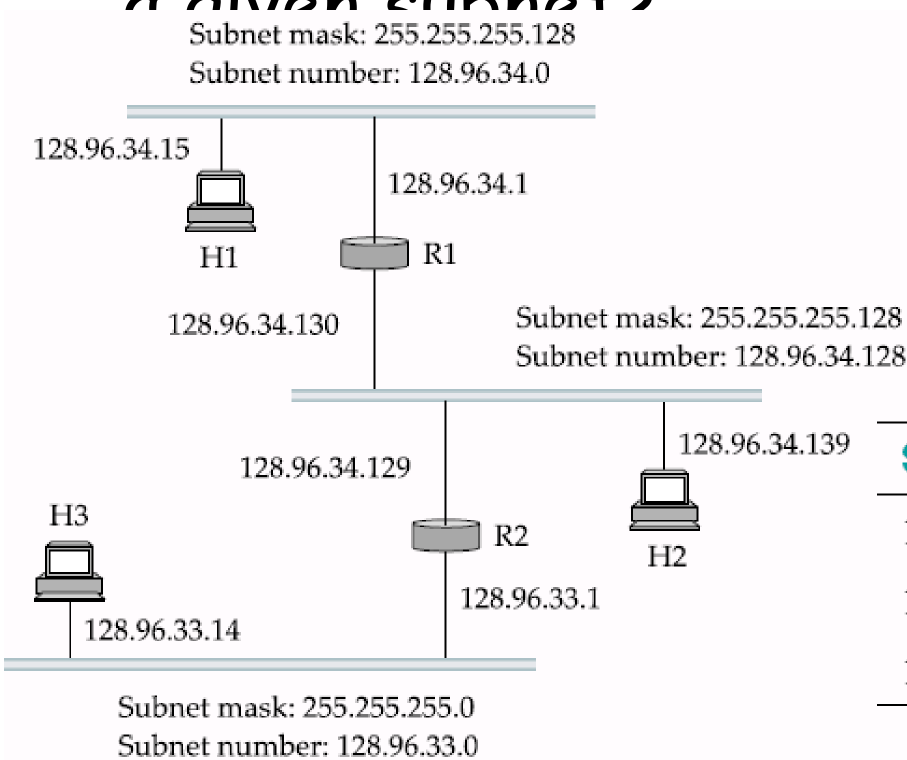

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Exercise C14b-1

- State to what next hop the IP packets addressed to each of the following destinations will be delivered

(a) 128.96.171.92
(b) 128.96.167.151
(c) 128.96.163.151
(d) 128.96.169.192
(e) 128.96.165.121

Table 3.19 Routing Table for Exercise 56

SubnetNumber	SubnetMask	NextHop
128.96.170.0	255.255.254.0	Interface 0
128.96.168.0	255.255.254.0	Interface 1
128.96.166.0	255.255.254.0	R2
128.96.164.0	255.255.252.0	R3
⟨default⟩		R4

Scaling Problem

- Need to address two scaling concerns in the Internet
  - The growth of backbone routing table as more and more network numbers need to be stored in them
  - Potential exhaustion of the 32-bit address space
- Address assignment efficiency
  - Arises because of the IP address structure with class A, B, and C addresses
  - Forces us to hand out network address space in fixed-size chunks of three very different sizes
    - A network with two hosts needs a class C address:
      - Address assignment efficiency = 2/255 = 0.78
    - A network with 256 hosts needs a class B address
      - Address assignment efficiency = 256/65535 = 0.39

First Attempt: No to Class B?

- Exhaustion of IP address space centers on exhaustion of the class B network numbers
- Solution
    - Say "NO" to any Autonomous System (AS) that requests a class B address unless they can show a need for something close to 64K addresses
    - Instead give them an appropriate number of class C addresses
    - For any AS with at least 256 hosts, we can guarantee an address space utilization of at least 50%
- What is the problem with this solution?

First Attempt: No to Class B?

- Problem with this solution
  - Excessive storage requirement at the routers.
- If a single AS has, say 16 class C network numbers assigned to it,
  - Every Internet backbone router needs 16 entries in its routing tables for that AS
  - This is true, even if the path to every one of these networks is the same
- If we had assigned a class B address to the AS
  - The same routing information can be stored in one entry
  - Efficiency = 16 × 255 / 65, 536 = 6.2%

Addressing Scaling Problem

- Classless Inter-Domain Routing (CIDR)
  - Addresses two scaling concerns in the Internet
    - Backbone routing tables are getting big
    - Potential exhaustion of the 32-bit address space
  - Balancing two factors
    - Minimize the number of routes that a router needs to know
    - Allocate addresses efficiently.
  - CIDR uses aggregate routes
    - Uses a single entry in the forwarding table to tell the router how to reach a lot of different networks
    - Breaks the rigid boundaries between address classes

Classless Addressing: Example

- Consider an AS with 16 class C network numbers.
- Instead of handing out 16 addresses at random, hand out a block of contiguous class C addresses
- Suppose we assign the class C network numbers from 192.4.16 through 192.4.31
- Observe that top 20 bits of all the addresses in this range are the same (11000000 00000100 0001)
  - We have created a 20-bit network number (which is in between class B network number and class C number)
- Requires to hand out blocks of class C addresses that share a common prefix (sometimes, called supnetting)

Classless Addressing: Notation

- Requires to hand out blocks of addresses that share a common prefix
- Convention
  - Place a /X after the prefix where X is the prefix length in bits
- Example
  - 20-bit prefix for all the networks 192.4.16 through 192.4.31: 192.4.16/20
  - A single class C network number, which is 24 bits long: 192.4.16/24

Routing and Classless Addressing

- How do the routing protocols handle this classless addresses
  - It must understand that the network number may be of any length
- Represent network number with a single pair

```
<length, value>
```

- All routers must understand CIDR addressing

Routing and Classless Addressing: Example

- Route aggregation with CIDR

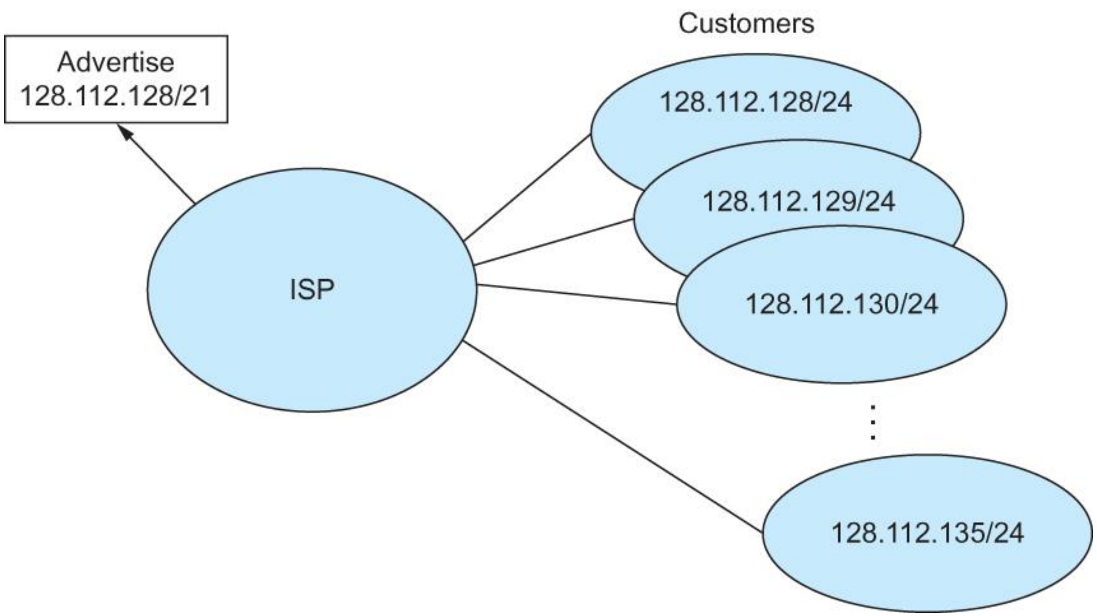

IP Forwarding Revisited

- Original *assumptions* in IP forwarding mechanism
  - It can find the network number from destination IP address in a packet
  - Then look up that number in the forwarding table
- Need to *change* this assumption in case of CIDR

IP Forwarding in CIDR

- Prefixes may be of any length, from 2 to 32 bits
  - Prefixes in the forwarding tables may overlap
- Some addresses may match more than one prefix
  - Example
    - Both 171.69/16 and 171.69.10/24 may coexist in the forwarding table of a single router
    - A packet destined to 171.69.10.5 clearly matches both prefixes.
  - Principle of "longest match"
    - A packet destined to 171.69.10.5 matches prefix 171.69.10/24
    - A packet destined to 171.69.20.5 matches 171.69/16

Exercise C14b-2

- State to what next hop the IP packets addressed to each of the following destinations will be delivered

(a) C4.4B.31.2E
(b) C4.5E.05.09
(c) C4.4D.31.2E
(d) C4.5E.03.87
(e) C4.5E.7F.12
(f) C4.5E.D1.02

Table 3.21 Routing Table for Exercise 73

Net/MaskLength	Nexthop
C4.5E.2.0/23	A
C4.5E.4.0/22	B
C4.5E.C0.0/19	C
C4.5E.40.0/18	D
C4.4C.0.0/14	E
C0.0.0.0/2	F
80.0.0.0/1	G

Questions?

- Subnetting
- Supernetting
- Classless addressing routing
    - Network prefix and length of prefix